LARGE BUSINESS MAINFRAMES

FIVE YEAR FORECAST

INPUT

Y-SBH
LBM

CUSTOM

AUTHOR

Large Business Mainframes 5yr

TITLE

DATE LOANED	BORROWER'S NAME

COVER AND BINDING BY 3M BRAND BINDING SYSTEMS

LARGE BUSINESS MAINFRAMES

FIVE YEAR FORECAST

FOR:

MR. PETER LABE
SMITH BARNEY, HARRIS UPHAM
1345 AVENUE OF THE AMERICAS
NEW YORK, NY 10019

MARCH 21, 1977

INPUT LIBRARY

INPUT

Digitized by the Internet Archive
in 2015

https://archive.org/details/largebusinessfivunse

TABLE OF CONTENTS

INTRODUCTION

THE CHANGING COMPETITIVE ENVIRONMENT

GROWTH OF THE LARGE MAINFRAME MARKET

BURROUGHS CORPORATION

CONTROL DATA CORPORATION

HONEYWELL

IBM's LARGE MAINFRAME PROCESSORS

SPERRY UNIVAC

APPENDIX - QUESTIONNAIRES

INTRODUCTION

- Over the next five years the share of the large business systems will increase at the expense of the medium size system.

- The new systems will be evolutionary not revolutionary in order to preserve the user's software investment.

- Wherever possible, firmware will replace software:
    - to lower the total cost of using the system.
    - to simplify the use of the system.
    - to increase the system's maintainability.

- There will be an acceleration toward:
    - communications-oriented processors.
    - software that is simpler to use.
    - terminals that incorporate instructional programs.
    - editing routines to help the user avoid errors.
    - specialized industry packages which are modular and can be modified by the user.

- The management of D.P. is increasingly reporting to the president or to the executive committee rather than to the treasurer, controller or the V.P. of finance; thus, the new applications being emphasized encompass the planning of the total business environment rather than conventional accounting applications.

- The study team evaluating large new business systems is a multidisciplinary group whose objective is to select equipment which will become

integral to the management of the business of the company. Since the time required to implement large complex systems is long (1 and 2 years), the installation life is projected to be an average of 7 years. Therefore, there is a significant increase in the purchase of large scale systems by the users. In 1976 over 75% of the large business systems installed were purchased.

- The motivation for replacing exisiting systems has to be:
  - A two-fold improvement in price/performance like that achieved by Amdahl. In addition, a system has to be provided which is either plug-compatible or easily convertible, eliminating the expense and the time involved in a conversion.
  - Insufficient capacity.
  - Inadequate service.
  - Desire to implement new applications not available on the old system.

- IBM strategies in the large business system area are to:
  - Preserve the customers' programming investment.
  - Increase number of applications by expanding into auxiliary areas such as word processing.
  - Improve price/performance.
  - Respond to competitive threats without impacting present revenues.
  - Continue to extend present 370 architecture by use of SNA and new peripherals toward integrated multifunctional multiprocessor network-oriented systems.

- The 370/158s and 168s are the most successful models in IBM's large scale business systems.

- There were over 2,000 IBM 370/155, 158, 165 and 168 systems installed at the end of 1976. Over 65% were purchased systems.

- IBM is delivering 15 370/168s per month. The backlog on the 360/168 is 8 to 12 months.

- In the large business system area, the customer is sophisticated and sensitive to price/performance advantages.

- Profit margins are higher at the upper end of the mainframe business spectrum.

- Despite the fact that Amdahl's inroads into IBM's domain have to date affected less than 2% of IBM's installed base, IBM will be forced to react to the Amdahl threat.

- IBM's reaction is likely to be on two levels:
  - a 50%-80% improvement in performance
  - a 15%-20% decrease in price.

- These improvements will:
  - provide more performance to accounts which need it
  - maintain an image of leadership.

- The current IBM shipping rate is 500 to 600 large scale systems per year with a purchase ratio of over 50%.

- There is increasing competition in the large business system market from plug-compatible mainframes, traditional mainframe manufacturers, and megaminis.

- Software for large business systems is very complex. Efficient operating systems, data base management systems, and communications systems are essential tools in successfully implementing a large computer system.

- When IBM released its 360 software, it was not fully debugged or properly documented. As the 360 evolved into the 370, IBM kept upgrading its software in a patchwork mode. The reason was that IBM found itself in a trap of its own making.
  - Users didn't want to convert to new systems and lose their investment in software.
  - The only way software could take advantage of new hardware features and still be compatible was through a series of patches. IBM is getting around this in three ways.
    - Cleaning up existing software. This will be mostly accomplished by 1978.
    - Attracting the new non-computer user by evolving a whole new series of terminal-oriented programs and providing users with on-line instructional programs.

- Providing users with customized, reloadable control storage firmware to help him better utilize the hardware.

- IBM will continue to gradually improve the price/performance of its large business system mainframes after a step function improvement in price performance during 1977.

- IBM will continue to introduce improved peripherals, especially disk drives.

- IBM will introduce its new architecture based on the 370/115 and 370/125 after all the other options are exhausted.

THE CHANGING COMPETITIVE ENVIRONMENT

- IBM is and will continue to be the dominant force in the large G.P. mainframe market.

- In the past, IBM has had little difficulty overpowering its traditional mainframe competitors. There is no reason to assume that IBM's dominance of the large business system market will not continue.

- As long as IBM dominates the large business mainframe market, an orderly oligopolistic improvement in price/performance caused by IBM's desire to expand the market while continuing to dominate it prevails.

- Short price/performance discontinuities occurred in the market when IBM announced new equipment giving users large performance increases while increasing prices only moderately. (See Exhibit

- IBM's moves are usually reactions to competitive market pressures designed to maintain or increase its market share. Thus, the 360 and the 370 were announced at the time when IBM's existing products were being displaced. The 370/138 and 370/148 are similar reactions.

- There are now three new competitive factors which are causing IBM to lose market share in the large business system market. They are:
    - Success by the plug-compatible CPU vendors.
    - The price/performance erosion caused by mega-minis.
    - Increasing market activity by the Japanese companies in the large business market.

- Of the three factors mentioned, the most significant is the plug-compatible CPU vendor. His presence in the market will not only accelerate the price/performance improvement in the large CPU market, as will the other two factors , but more importantly, he gives users "the option to choose between an IBM mainframe and one that acts just like it."

- Large EDP systems have long payout records. The plug-to-plug CPU companies are designing a functional equivalent of 370 processors utilizing the latest semiconductor technology and pricing it below IBM levels. These companies have an appeal to a growing minority of the IBM user base.

- Amdahl is the major successful example of this approach, but other companies, such as Itel and CDC, have also announced their own compatible systems.

- Suddenly IBM becomes an industry standard by virtue of its market dominance and its imitators.

- Thus, by offering IBM plug-compatible mainframes, Amdahl, Itel and CDC may in fact be strengthening IBM's long term competitive position, since the defectors feel secure that they will always have the option to return to the new IBM mainframes to be introduced in the 1980s.

- Who is then getting hurt? No one yet — however, in the long run, it may be that some of the other large mainframe manufacturers will.

- The vise of the plug-compatible mainframe manufacturers leaves the other mainframe manufacturers in a dilemma.

- If they acquire IBM compatibility, their cost for software development is reduced, but they lose their uniqueness.

- If they preserve their unique software and architecture, the future pressure to improve price/performance will reduce their margins to the point where they will increasingly have difficulty offering a price/performance advantage over IBM or maintaining a sufficiently high level of R&D to continue competing in the market.

- Because of these factors, the impact of the plug compatible mainframes on the market structure may be much more insidious than is readily apparent.

- The megamini market alternative is becoming a very significant challenge to IBM, primarily at the bottom and middle range of the 370 market. The major minicomputer suppliers, such as Digital Equipment, Data General, and Hewlett Packard, are demonstrating DP equivalents at 30%-60% lower purchase prices than IBM and other traditional EDP system suppliers. This threat is emphasized by studies aimed at replacing large mainframes with networks of minis (e.g., Citibank).

- The ability of minicomputer companies to offer lower prices is caused by their lower profit margin targets, purchase-only pricing and their lack of need to protect a rental base.

- IBM will become more price/performance competitive with the megamini alternative or face the prospect of losing market share.

- By introducing the Series 1 mini, IBM has shown that it is willing to face competition on a targeted, limited-response basis in order to protect its traditional markets. This response will certainly be expanded in the future.

- Japanese Computer Companies. Companies whose computer business is controlled by the government, such as Fujitsu and Hitachi, may have priorities other than profitability. These companies are responsive to social and political goals in exchange for government support of their R & D effort.

- Japan's government EDP plan calls for a share of the export markets from the current 2% to about 15% in the 1980s.

- There is no doubt that this will be at the expense of the U.S. companies who currently dominate the overseas markets.

- Responding to the Challenges . The manufacturers of large business systems facing these threats must concentrate on three strategies.
    - Decide what price/performance improvement is going to be required to keep them competitive.
    - Decide what level of hardware & software R&D expenditures will be needed to keep them in the game.
    - Determine whether their user base is large enough to allow them to support that level of R&D and still maintain a reasonable profit level.

- Each manufacturer of LBM systems is aware of these challenges, and each is confident that he can overcome them. We will analyze each vendor's approach and liklihood of success individually after examining the condition of the large mainframe market.

GROWTH OF THE LARGE MAINFRAME MARKET

- The installed base in the large mainframe market is estimated to have grown from $2.3 billion in 1971 to $7.4 billion in 1976. The share of market of the large mainframes grew from 6.4% to 11.8% of the total market during the same time period.

- The large mainframes will continue to grow over the next five years, increasing their market share from 11.8% to 14.1% of the installed base. (See Exhibit I.)

- IBM is estimated to derive between 38% and 40% of its revenues on an if-sold basis from their large systems. These include 70% systems in the $2 to $4 million category and 30% systems in the $5 to $8 million category. (See Exhibit II.)

- Univac is estimated to derive 49% to 53% of its revenues from the large systems. Univac's large systems, ranging from $2 to $2.5 million, account for 40% of total large business systems revenues, while the very large, ranging from $3.4 to $5 million, account for 60%.

- Burroughs derives only 20% of its revenues from the large class of business processors. Its large processors account for 85% of the large business processor revenues and average around $3 million on an "if sold" basis. The very large processors, accounting for 15% of the large processor revenues or only 3% of corporate revenues, average over $6 million on an "if sold" basis.

EXHIBIT I

SHARE OF MARKET

OF LARGE MAINFRAMES,

INSTALLED BASE

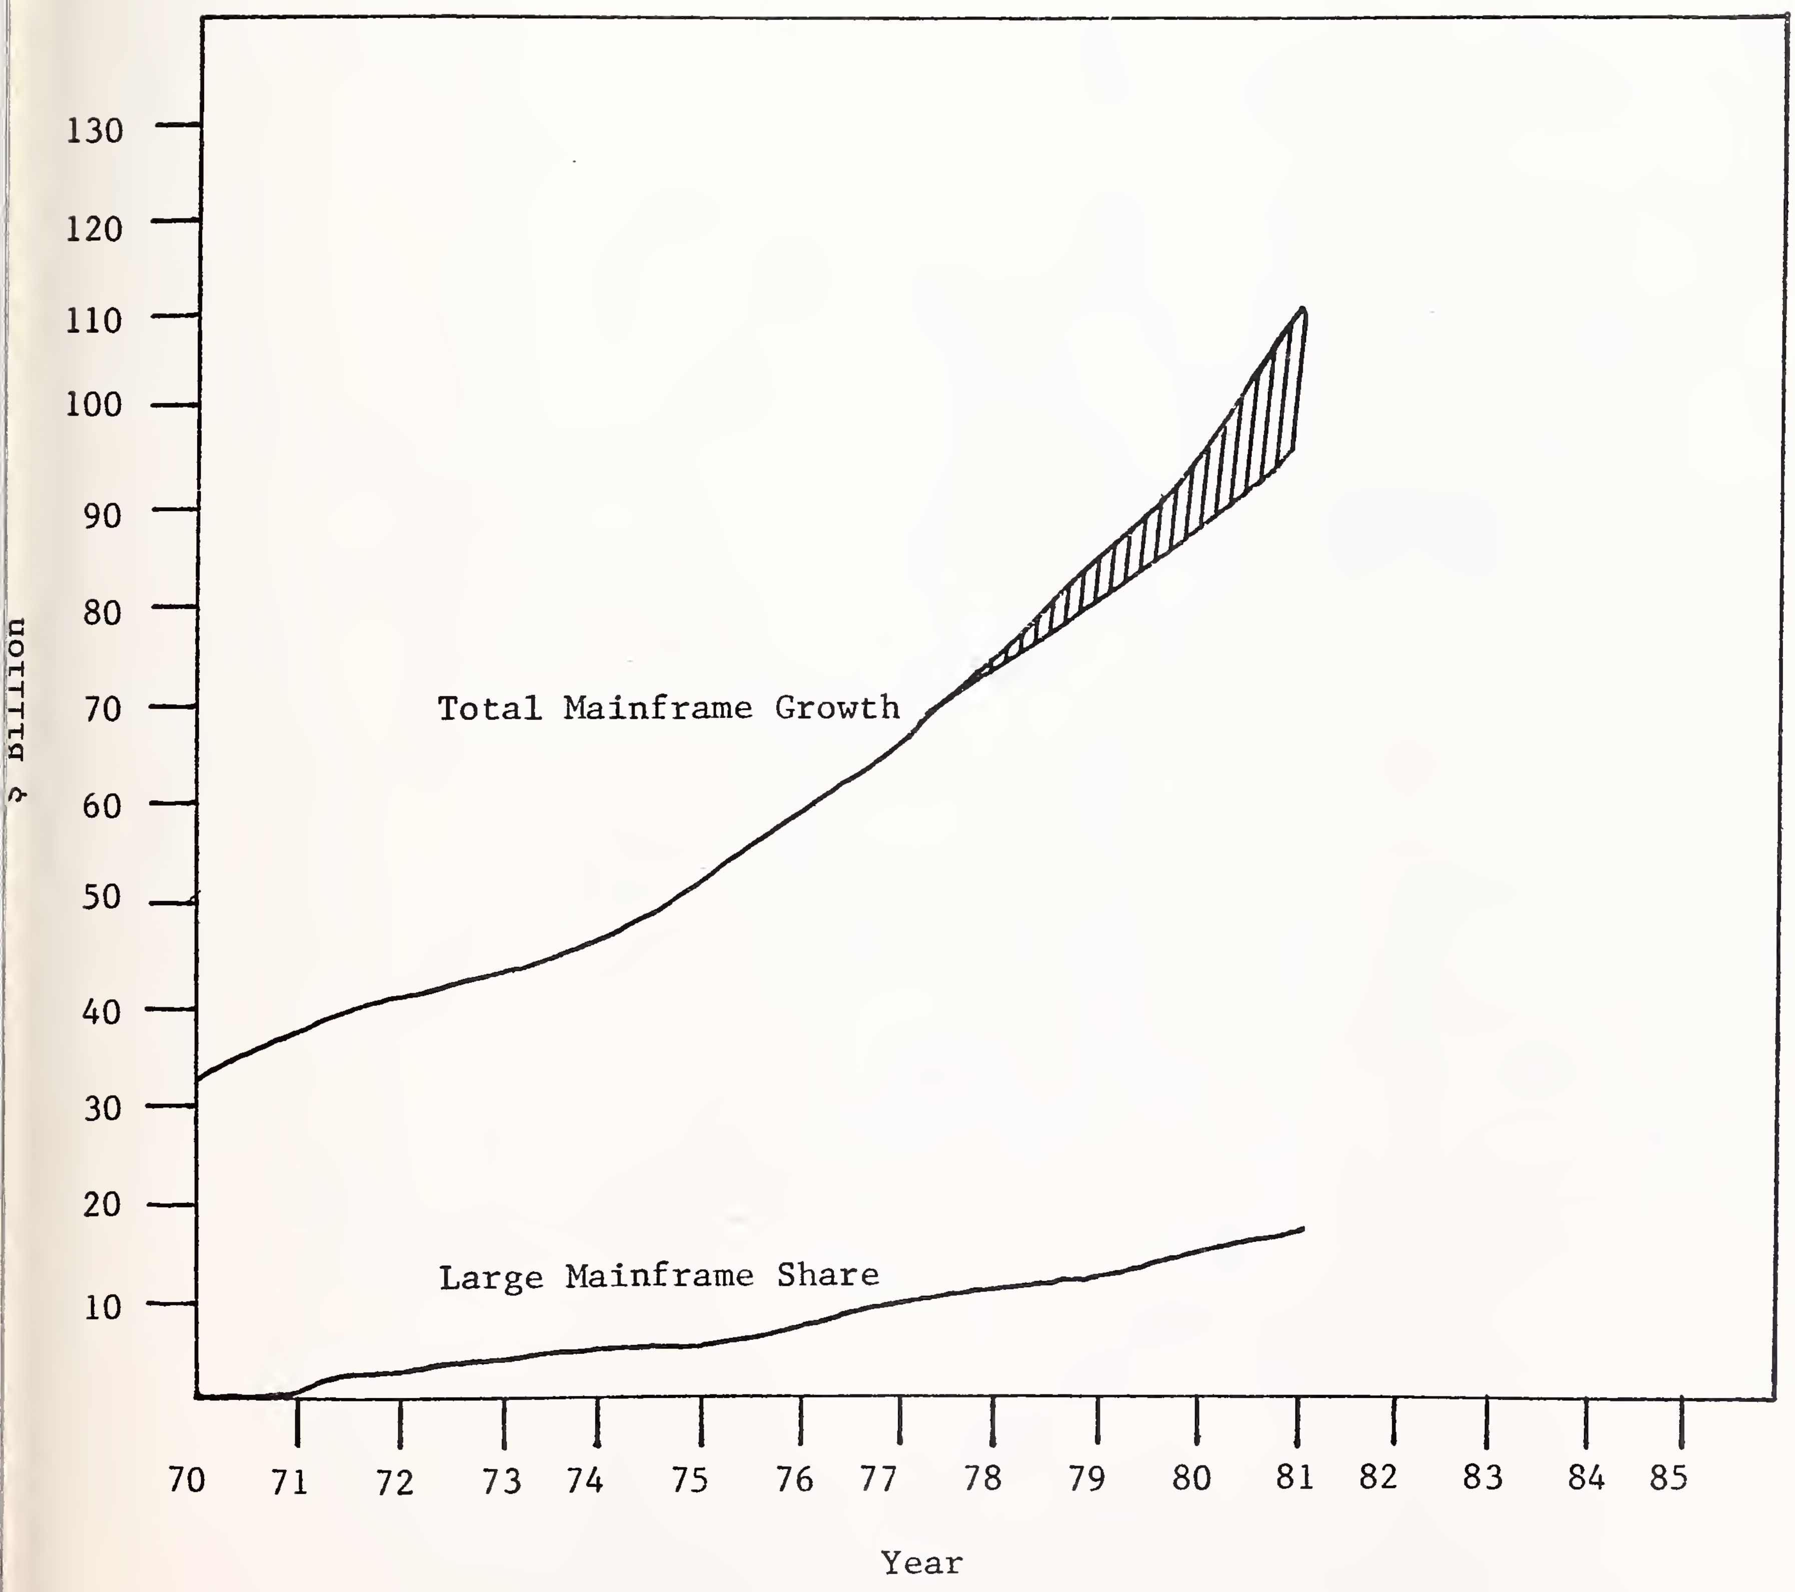

EXHIBIT II

1976 SHARE OF INSTALLED LBS BY COMPANY

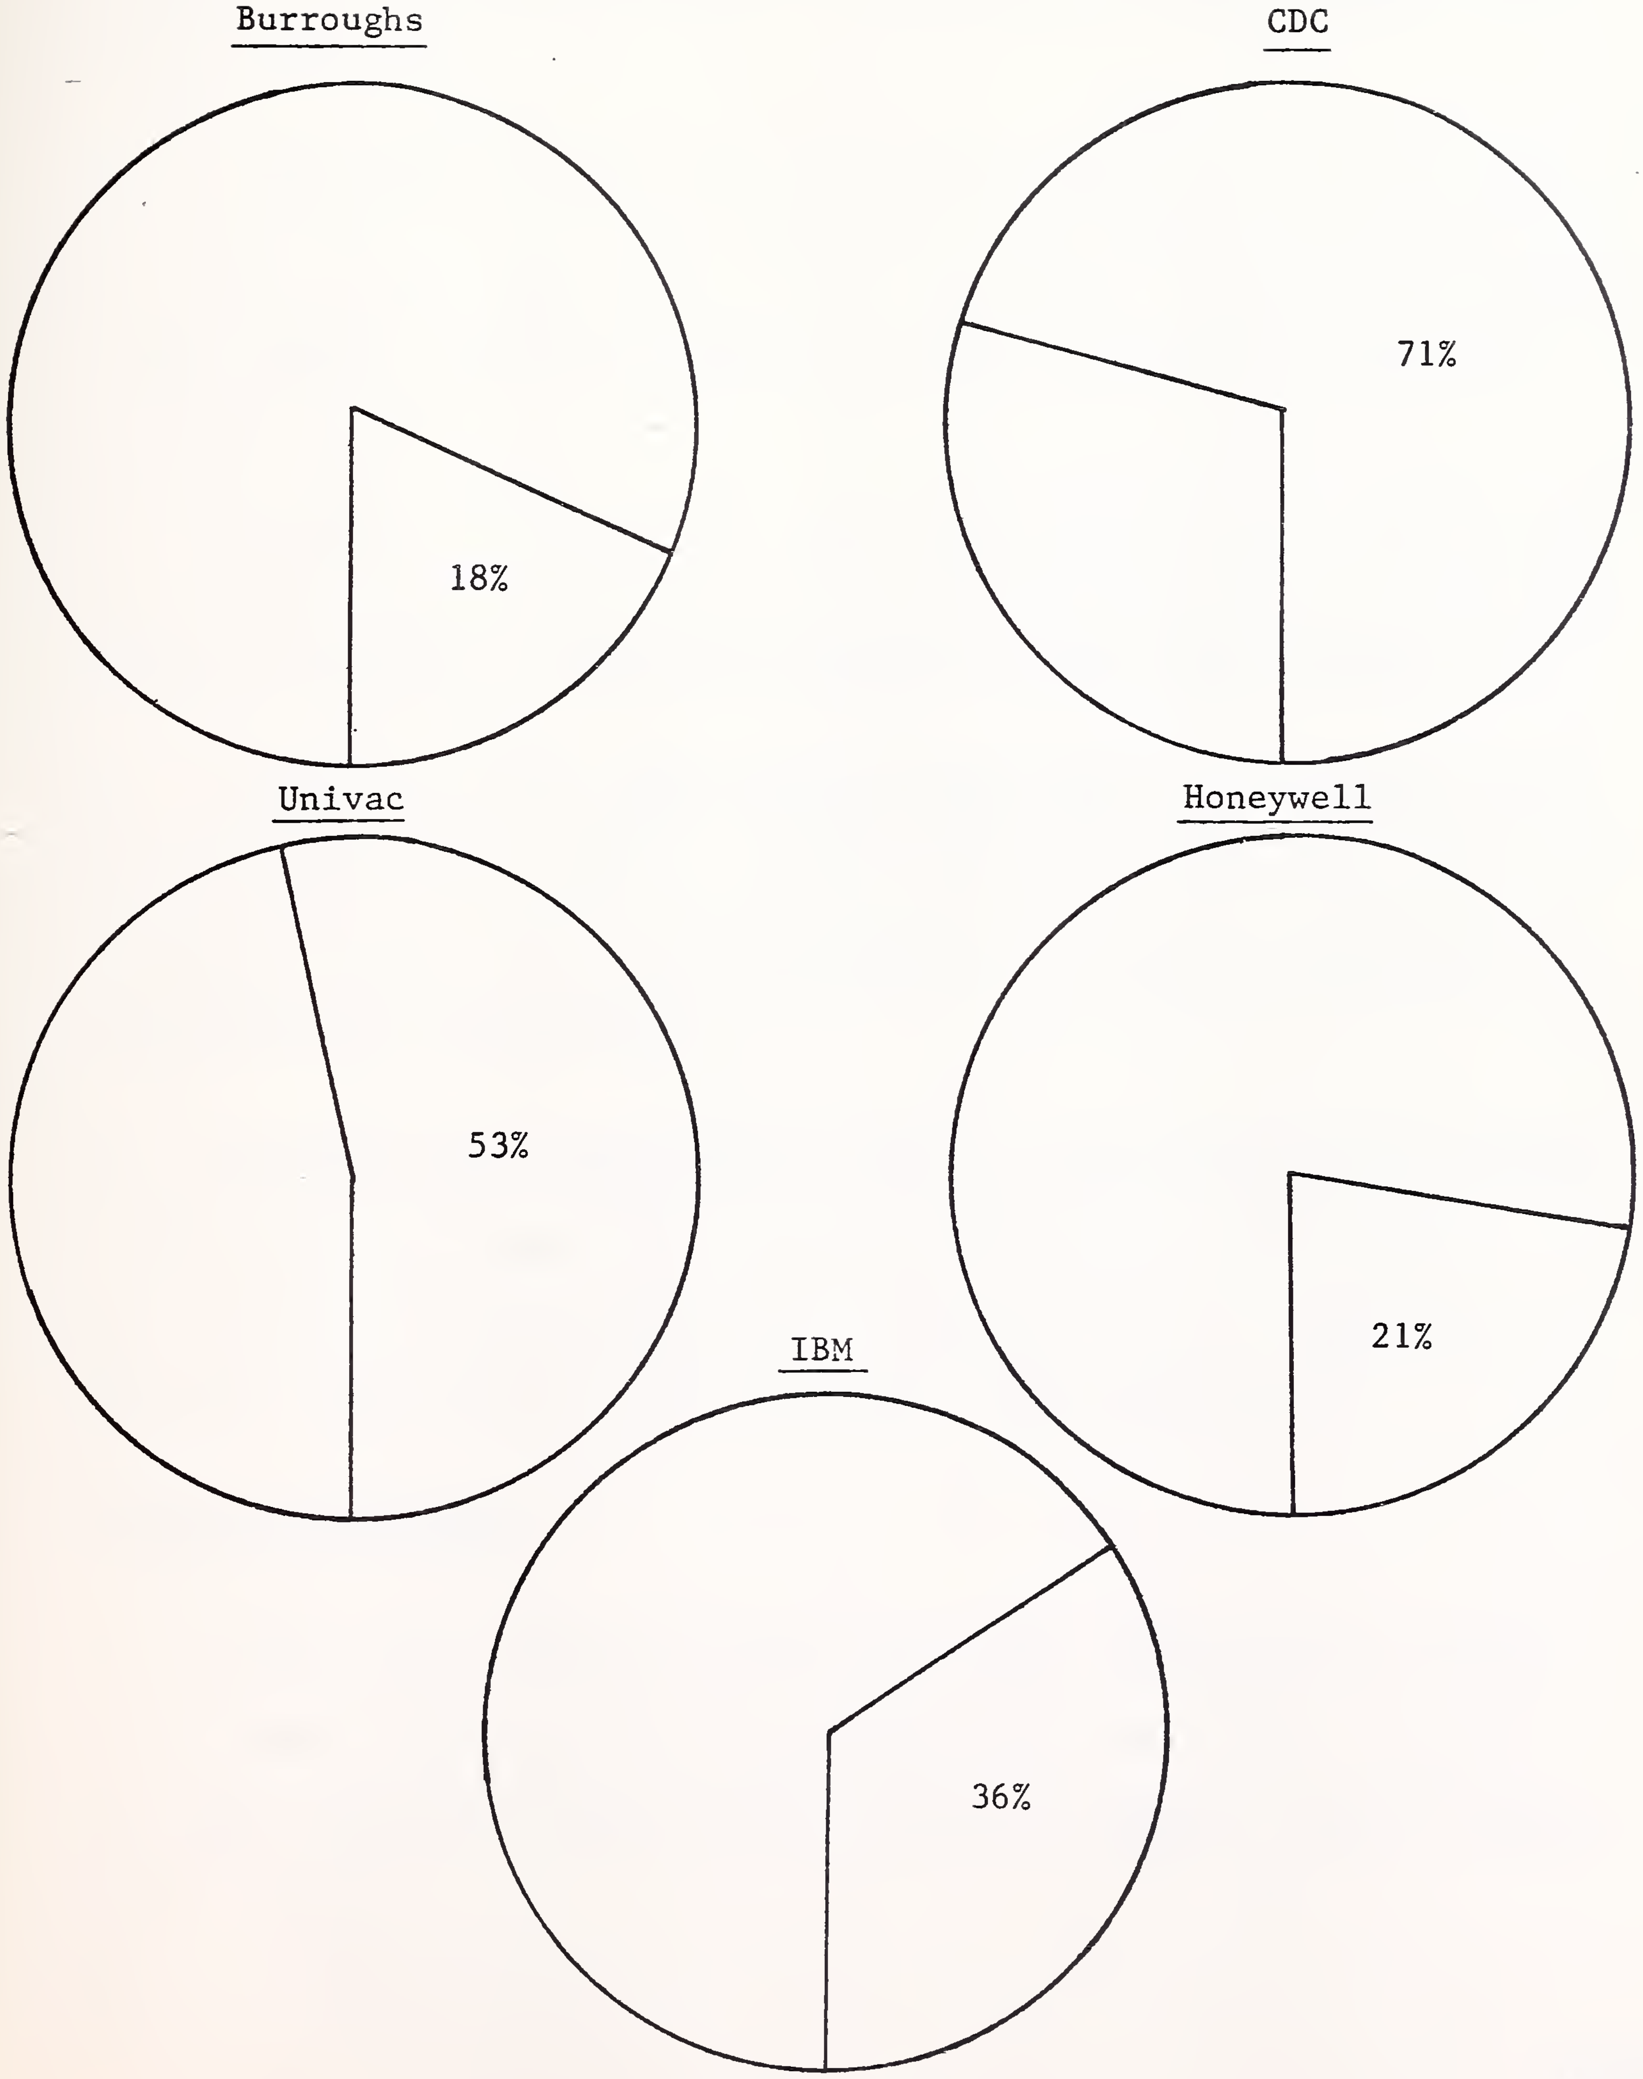

- Honeywell derives between 21% and 25% of its revenues from large business systems. Only 7% of Honeywell's systems are in the very large category, averaging $4.5 million. 93% are in the large category, averaging $2.8 million.

- 70% of CDC's revenues are derived from the large mainframe market, 54% from the very large, and 46% from the large. The systems in the very large category average $6 million and those in the large average $2.5 million.

- IBM has its systems evenly distributed along industry lines; thus, IBM has over 80% of its systems installed in eight major markets.

- CDC, on the other hand, derives almost 100% of its revenues from the same eight major markets. In fact, CDC derives more than 75% of its hardware revenues from three major markets: the federal government, manufacturing, and education.

- Burroughs has a preponderant number of its installations in banks, while Univac has almost 40% of its installations with the federal, the state, and the local government. Honeywell has the highest proportion of its systems in manufacturing. (See Exhibit III)

TECHNOLOGY FORECASTS

- Despite the temporary breakthroughs other companies may achieve, it is IBM's technology and pricing which will, in the long run, have the most significant impact on the rest of the industry. This is due to three factors.

EXHIBIT III

INSTALLED BASE DISTRIBUTION BY COMPANY-
BY 8 MAJOR INDUSTRY MARKETS

INDUSTRY	%	IBM	HIS	UNIVAC	BURROUGHS	CDC
Business Services		14.6	13.2	18.8	18.2	20.9
Federal Government		6.1	14.8	18.0	8.7	29.0
Banking		7.6	4.4	1.1	22.8	0
Insurance		8.2	5.2	4.1	.98	.08
Education		5.4	3.5	8.4	5.6	18.6
State and Local Government		5.7	5.3	9.4	8.0	1.6
Wholesale		4.4	5.2	2.8	3.9	3.0
Nonelectric Machine Manufacturer		3.5	5.2	18.0	4.7	5.4
Other Manufacturers		23.5	26.0	18.0	17.6	20.0

- IBM dominates the large business system market.
- IBM's R&D budget exceeds $1 billion per year.
- All other companies peg their prices at some fraction of IBM's prices.

- The distinction between a general purpose market and the scientific market is rapidly becoming blurred, although CDC's large systems are typically considered to be scientifically-oriented. The CDC customers are demanding more transaction processing as well as data base access and manipulation capability. CDC will provide both and probably some special decimal instructions as well in its announcement scheduled for Mar. 29.

- Burroughs has announced its scientific processor which will have a business oriented front end.

- IBM has upgraded its 3838 array processor by adding seven new vector arithmetic algorithms under multiple virtual storage (MVS). In Fact, MVS is the key to the evolving IBM computer structure.

- This new structure is based on multiple functional processors whose functions may be altered through reloadable control store (RCS) to conform to the user requirements; thus, a processor may be performing large scientific calculations during the night and interactive data base and communication oriented transactions during the day. (See Appendix.)

- Hardware will be an increasingly smaller concern of users, and software flexibility and system availability will become their major concerns.

EXHIBIT IV

IBM KOPS/$ LBS

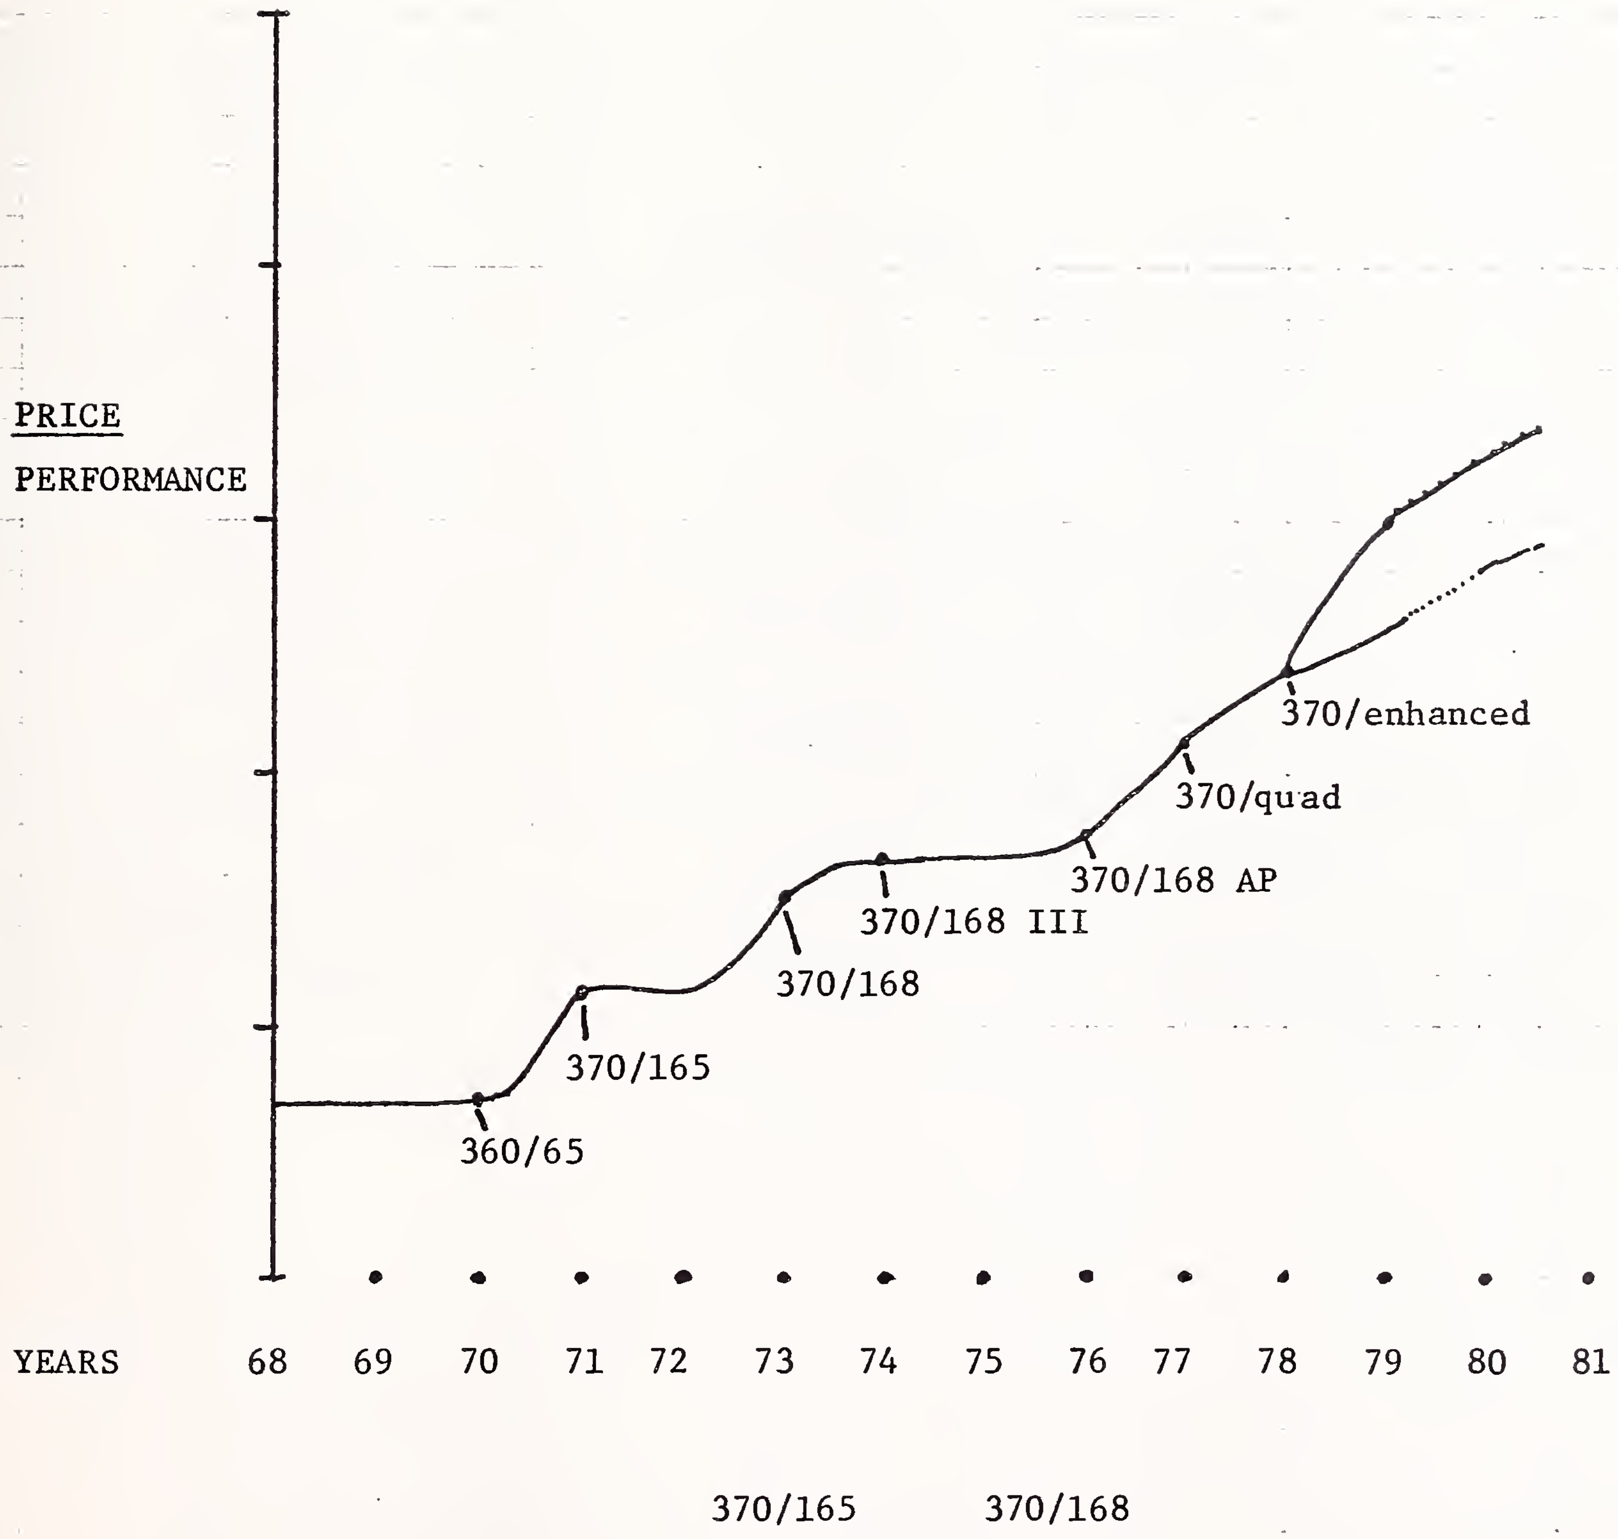

EXHIBIT IV (Cont'd.)

	MODEL	PURCHASE PRICE	KOPS	MEAN KOPS/ DOLLAR
1968	360/65	2400-2600	800-900	.360
1971	370/165	2200-2400	1300-1400	.590
1973	370/168	2900-3200	2100-2400	.750
1974	370/168-3	3200-3400	2500-2700	.810
1976	370/168 AP	4700-4900	3900-4200	.840
1977	370/178	3400-3800	3900-4300	1.150

EXHIBIT V

LBS SHIPMENT FORECAST

	1976	1981
IBM	2,160	4.000 - 4,500
Burroughs	150	400 - 480
Honeywell	231	360 - 440
CDC	248	340 - 400
Sperry	583	1000 - 1,200
Amdahl	92	400 - 500

$ MILLIONS

BURROUGHS CORPORATION

- On February 1, 1977, Burroughs announced a new series of large scale computers designed for large scale data base management systems and large on-line networks.

- The 7800 scheduled for delivery in the first quarter of 1978 has two models:
  - The 7811 - a single CPU and a single I/O processor.
  - The 7821 - a dual CPU and a dual I/O processor.

- Price of the 7811 with 1.6 M. characters is $2.5 million, while the dual processor with 6.3 M. characters sells at $4.4 million. This puts Burroughs squarely in the 370/158 and 370/168 market.

- This is the latest of the large scale Burroughs systems and is compatible with the previously announced 6800, 6700, and 7700 series. Like its predecessors, it uses a fixed length data word of 48 bits plus hour control and priority bits.

- The 6700 series has four models as does the 7700. The CPU prices range from $334 thousand for the 6738 model to $4.6 million for the 7785, a four processor 210P system.

- The key to the Burroughs large business systems is the Master Control Program (MCP), the Burroughs operating system designed to perform in the multiprogramming and multiprocessing environment. This is the most sophisticated operating system in the industry, having used virtual memory techniques six years before IBM endorsed the concept. In addition, MCP uses stack processing and data descriptors for memory allo-

cation and, and it provides for a fail-soft method of handling changes in resource availability.

- Burroughs provides time-sharing, remote batch, message switching, and a terminal management system for its large business data processing systems. In addition, Burroughs has introduced the Burroughs Data Link Control (BDLC), a flexible bit-oriented communication protocol.

- The Data Base Management system provided by Burroughs is designed to operate either on-line or on a batch basis. The user has the option of organizing his files into a hierarchical, a network, or a ring structure.

- All the Burroughs systems, including the new ones, are still built on the channel-to-main-memory building block basis. This type of architecture consists of one or more separate central processors, memories, or I/O processors. Others using this type of architecture are Univac, with the 1100 series, Honeywell, with Series 60 Level, 66, and 68, and the Amdahl 470V/6 in the large business market. This type of architecture permits easy upgrades by the addition of more CPUs or I/O processors.

- Burroughs now offers both core and IC memories on its large business systems. Low voltage, low current bipolar circuitry manufactured at its Rancho Bernardo plant, and the Current Mode Logic (CML), have been incorporated into the 7800 systems, enabling the company to improve performance approximately 2.5 times over its predecessor at a 30% increase in price.

- User programs for the Burroughs LBS are all written in higher level languages such as Fortran, Basic, Algol, APL, RPG, PL/1 and Cobol.

- Burroughs does not provide hardware emulation for competitive systems; however, it does offer software conversion aids enabling 360/370 Cobol to be translated to the 6800/7800 Cobol.

- Burroughs PL/1 is compatible with IBM's version of the same language, while BDLC and IBM's SDLC are sufficiently similar to allow the two networks to communicate with each other.

- Burroughs has announced that it is developing a new scientific super computer. This system aimed at CDC and CRAY can use a B7700/7800 as its system manager, providing it with a balanced mix of capabilities including time-sharing and general purpose computing.

- There is no reason to assume that a processor emulating the 370 could not be inserted into this architecture.

- Burroughs large business systems trace their ancestry to 1961 when the B5000 was first announced. Most of the advanced features of that system have been preserved in the follow-on versions.

- The high end of the Burroughs 6000 series and the Burroughs 7000 series compete with the IBM 158 and 168 MP as well as with the Univac 1100 series, Honeywell's Level 66 Models 60, 80 and 85, and CDC's Cyber 173, 174 and 175.

- Burroughs competes favorably against IBM in the multiprocessor mode, because the MCP was originally designed as a multiprocessing operating system and has been in operation on the 6700 machines since 1969.

- Burroughs has made its major penetration into the domestic market with banks and federal, state and local government agencies.

- Burroughs LBS have achieved a growing market share. In fact over 40% of the Burroughs systems are replacements of other manufacturers, while 40% are upgrades from smaller Burroughs systems.

- Management's View of User Requirements. Burroughs management perceives the major user concerns to be as follows:

  - Improvement in price/performance.
  - Easy use of large data bases.
  - Increased reliability, availability, mantainability.
  - Ability to communicate effectively.
  - Desire to preserve software investment.

- It feels that Burroughs' recent announcements have been aimed at providing users with what they have been asking for. Burroughs believes it has the highest user loyalty of any manufacturer.

- The component technology Burroughs has developed gives the company a significant proprietary edge. The software and systems architecture are, however, the functions most appreciated by the Burroughs users.

- User Reactions. Of the 7 Burroughs users interviewed, 6 were very happy with their system. One had had a memory replaced. Unanimously, they praised the Burroughs software and its ease of use.

  - 3 of the 7 thought Burroughs maintenance was "bad" or "lousy."
  - 1 user thought Burroughs maintenance improved in the last 18 months. (They are finally getting their act together.)

- 1 user thought Burroughs marketing was "atrocious."
- All users will be either upgrading mainframes or adding communications processors or terminals.
- 5 of the 7 will be attaching minis on-line to the mainframes over the next 5 years. Most aren't sure how many. They all feel that the work load of the host will increase with the addition of the on-line minis.

● Although most users interviewed felt that minis would be integrated into their systems, most were still uncertain about how this would be done and who would be responsible for what.

● 2 of the 7 large Burroughs mainframe users interviewed were previous IBM users. The others were upgrades.

- Most Burroughs mainframe users have the mainframe for 5 years before upgrading. In a multiprocessor environment, upgrades can be done modularly without disrupting the total system's operation.
- 4 of the 7 users have no concern about being obsoleted.
- None of the 7 is concerned about compatibility with IBM.
- The following are reasons cited for changing mainframe suppliers:
    - Inability to support the system.
    - New software development.
    - Dramatic improvement in price/performance (2) 50% or more.
    - Technological breakthroughs.
- Burroughs was selected because of :
    - ease of conversion (3)
    - better price/performance (2)
    - satisfactory benchmarks (1)
    - "best deal" (1).

- It is interesting to note that no one mentioned that software was the best-liked feature after the system was installed.

- Conclusion. Burroughs has taken an extremely aggressive stance in the LBS market. The price performance of its systems is quite competitive, and, depending on the configuration, it has 25%-30% better price/performance than the equivalent IBM large business system.

- Burroughs has been increasing its share of the LBS market going from 3.7% in 1971 to 5.3% in 1975. This expansion will continue to increase, because Burroughs has a solid base of medium size users who will be upgrading to larger systems in the next 5 years.

CONTROL DATA CORPORATION

- CDC would not have been interviewed except for the fact that it is getting ready to introduce a 370/148 compatible and possibly a 370/158 compatible system.

- These systems , originally designed by Cambridge Memories, are now going through final testing. The group within CDC responsible for these systems is the peripherals group the same group that manufactures and sells IBM plug-compatible peripherals.

- CDC doesn't consider the peripherals groups activities to be in the mainstream of its business, but rather an opportunistic thrust into a potentially lucrative market.

- CDC is best known for its large and very large scientific systems, such as the Cyber 26 and the Star 100.

- These systems are designed to solve very large computation problems. The evolution of their software is enabling them to fit into a network environment. These systems have no peer when it comes to performing iterative computations similar to those found in mathematical modeling.

- The Cyber 76 has a clock rate of 27.5 ns, as compared to Amdahl's 32 ns, TI's ASC 60 ns and CRAY-1 is 12.5 ns. The Cyber uses a 60-bit word as compared to Univac's and Honeywell's 36-bit words on the 1100 and 66 and 68 levels, respectively.

- The Cyber 170 line, which has a distributed processing architecture, is also oriented to solving large scientific problems. The computers

employ very fast emitter coupler logic and MOS memories.

- CDC perceives the major user concern today to be "a groundswell of uneasiness about the possibility of being stuck with a dinosaur, "while" the rest of the world is going to mini's."

- CDC has taken an active role in developing computer networks. The distributed architecture of its systems lends itself to the implementation of the network concept.

- CDC perceives the major user concern today (its users tend to be more sophisticated) to be "a groundswell of uneasiness about the possibility of being stuck with a dinosaur," while "the rest of the world is going to minis."

- Data management is a concern, and the development of a data base processor is a possibility. CDC believes that the data base software is still insufficiently well-defined to justify building a data base processor.

- Although CDC has had an excellent year of shipping large mainframes, it foresees a slowdown of growth in this market. Currently 71% of CDC's mainframe installation base consists of large mainframes. This ratio will change as CDC expands its product line both upward and downward. Shipments of the recently-announced CDC small computer system and its to be announced medium computer will absorb the slack caused by slower large mainframe shipments.

- CDC sees no breakthroughs in component technology. All its chips are developed in-house and are either built in-house or subcontracted. A bubble memory is being worked on.

- CDC believes that IBM will use an evolutionary approach in its new systems using firmware controlled functional processors enabling a change in function to adapt to the changing requirements in the job stream.

- CDC has dedicated sales forces selling to various markets, and it has been very successful selling to the education, seismic, atomic energy, and computer services markets as well as to markets overseas.

- The company is now concentrating its R&D on software to link its data management and data communications software. This is a response to increasing demands by the scientific community for interactive processing, information retrieval and transaction processing.

HONEYWELL

- The Series 60 is Honeywell's attempt to integrate the Honeywell, GE, Olivetti, and Bull product lines into a single upward compatible product line.

- The Series 60 has various levels, the lowest of which is Level 61 or entry level. The next is Level 62, then 64 and finally Levels 67 and 68, which represent the upper end of the product line.

- Within each level there are also a number of models.
  - Level 66 comprises Models 10, 20, 40, 60 and 80.
  - Level 68 comprises Models 60 and 80 as well as the recently announced Model 85.

- Honeywell's large business systems have their origin in the GE 600 series which in turn is based on military computer, the M60. The series 60 Levels 66 and 68 are extensions of the 600/6000 systems. These 36-bit word systems were designed primarily for commercial and transaction processing applications.

- A mature data base management system and a front end communications processor make the upper end of the Series 60 an excellent performer in the large distributed systems environment. The Multics software, developed jointly be Bell Labs, GE, MIT and the Government, allows the large multiprocessor Honeywell systems to function in a secure computer utility environment.

- The large sytems use Honeywell-designed LSI and 4K MOS memories. They are modularly designed, allowing extensive use of the distributed processing concepts. These large mainframes use a memory-oriented

structure with one to four system controllers controlling traffic between various system components and memory. The microprogrammed I/O multiplexer interferes with the communication processors and peripheral processors while concurrently executing programs.

- Honeywell offers virtual memory only on the Level 68 processors. The 68/80 includes:
  - A 2K word cache memory.
  - A decimal unit for business oriented instructions.
  - An associative memory for translating virtual addresses to real memory addresses.
  - An interrupt store "and" restore function to facilitate context switching.
  - Hardware to implement system security measures.

- In addition, the GCOS operating system has proven diagnostic routines designed to support a "fail soft" mode of operation on systems with multiple CPUs, I/O multiplexers and system control units.

- This fail soft capability is achieved by the use of the on-line testing system which performs diagnostic tests on each system module while the system is doing normal processing. The Honeywell error analysis and logging system heals, analyzes and logs errors associated with various systems components. If an unrecoverable error occurs, the component is bypassed, and the system continues to function.

- The Multics software and the distributed architecture of Levels 66 and 68 give Honeywell a potentially strong competitive position among:

  - Large companies which plan to implement their own in-house distributed networks.
  - Government agencies concerned with security.
  - Computer services companies.

- Honeywell tends to compete in selected situations, and it wins because it offers uniqueness.

Honeywell believes that ease of use is the most important single concern of users and that improving programmers' productivity is a close second.

- The component tecbnology development at Honeywell is behind schedule, and the newest Honeywell large business system, the 66/85, has liquid cooling to keep its circuits from overheating.

- The future increases in speed will be achieved mainly through architectural improvements, instruction segmentatation, parallel operations, macrofunctions, vectoring, or look-ahead -- nothing new, just taking concepts that have been around for years and making them work.

- Distributed processing is going through birth pangs, and it will go through 3 to 5 years of trauma before users can confidently install it.

- Honeywell sees a two-step function in price performance - one occurring now, paced by IBM, and another in 4 or 5 years, when some of the software currently being implemented is amortized and will be sufficiently stable to justify functional processors built around it.

- The Honeywell data base management language seems to be one of the more advanced on the market.

- Unlike Univac, which has managed to successfully consolidate its hardware into two major product lines and a small business computer, Honeywell is still struggling with a multiplicity of hardware and software systems.

- For example, the GE 635 software and GECOS are no longer supported, and the new software developed for the 66 and 68 series can't be run on the 635. Multics also seems to be in the process of being gradually replaced by a new version of GCOS which, however, will lack some of the Multics features. This will leave Honeywell exposed to competitive thrusts from both Univac and Burroughs - both of which are rapidly upgrading their product lines.

- Users' reactions to their Honeywell systems were unanimously very good. The system is very solid and easy to use. Conversion was easy for most users, since Honeywell provided the necessary manpower to help in the conversion - especially from the 8200 to the 60 series.

- The interviewed Honeywell users plan to keep their systems for 6 to 7 years, generally 1 to 2 years longer than the users of other vendors' systems.

- None of the Honeywell users is concerned with technical obsolescence or with IBM compatibility, because nonstandard clauses allow users to upgrade equipment at low prices.

- Honeywell has sold some of its amortized equipment to universities at deep discounts.

- The 2 complaints of Honeywell users are that software documentation is slow and that peripherals are the weakest point in an otherwise extremely reliable system.

- The interviewed users seem to have been captured as a result of a good understanding of the user's problem and the ability to convince him that the proposed equipment will best solve his problem. Good salesmanship was mentioned by 3 users in addition to benchmarks, which were mentioned by 5. In all instances the competition was intense.

- Over the next 5 years all the users see more communications gear and terminals being added to the system. One plans to double the size of installation, and one plans to increase use of data bases and on-line inquiry capability. The addition of processors is as much a function of additional capacity requirements as it is a function of a "fail soft" requirement.

- The interviewed Honeywell users indicated either a disinterest or an unawareness of "distributed processing". However, they all utilize on-line terminals and RJE terminals to enter data into the larger system.

- Of the 7 users interviewed, 3 selected Honeywell because of price, 3 because they had run the benchmarks, and 1 because of the support offered. In each case the salesman knew what the user wanted.

CONCLUSION

- Honeywell is structurally the weakest of the mainframe competitors. It has still not fully integrated its product line. The Distributed Processing thrust coupled with the very solid hardware and software could provide a growth opportunity if costs are kept in line.

IBM's LARGE MAINFRAME PROCESSORS

- IBM's current large processors (the 370/158 and 370/168) are the upper end of the 370 product line. Since its announcement of the System/370 in 1970, IBM has enhanced its capabilities by adding virtual memory and multiprocessing capabilites as well as a host of new peripherals and a new network management protocol.

- The recent additons are the 158-3, the 168-3 and the attached processor system; each provides incremental increases in price/performance over its immediate predecessors.

- When IBM first announced its 370/155 and 370/165, it obtained a doubling of price performance through the use of the "cache" memory. In every other respect technology was the same, and programs were compatible with the 360 series, making conversion very easy.

- Virtual storage capability was added in 1972, and dual processors were announced in 1973, while main memory was increased from 4 to 8 million bytes.

- In 1976 the attached processor system, consisting of a second instruction processing unit with its own 32K high speed buffer, was announced, providing a 1.5 to 1.8 improvement in internal performance at a 15% to 20% increase in price. A.P.S. didn't slow Amdahl down significantly however, it did provide IBM users an upward growth path.

- Until 1973 the IBM system user didn't have the option of a multiprocessor configuration. When the MP configuration was announced, 2 basically different approaches were made available.

- The "tightly coupled" system in which 2 large processors share the same combined main memory and operate under a single VS system.
- The "loosely coupled" system which can have up to 4 individual processors operating under its own individual operating system using its own main memory, but controlled by an overall "systems manager" and a common job entry subsystem.

- The major complaint of IBM users is the high cost of software overhead needed to manage the IBM multiprocessing system. Competitors like Burroughs, Univac, Honeywell and CDC have had MP systems and MP software as major selling points for a long time. Their system overhead for running a multiprocessor system has now been reduced to 10% to 15% while IBM's overhead on some jobs can still run as high as 30%.

- The most important technological development of the 370 line is the reloadable control storage (RCS), a microprogrammable semiconductor memory which controls the internal operations of models 370/115 through 370/158. The instructions are loaded into the RCS through a floppy disk. Different programs, features, or even new instructions can be inserted into the RCS, effectively reconfiguring the system or adding features such as floating point arithmetic, emulators, language processors or microsequences.

- Although RCS isn't presently available on the 168, a reconfigured architecture could provide the new IBM large business systems with microprogramming capability.

- IBM is expected to announce an upgraded 168 system within the next month. We expect that this system will stay within the parameters outlined in the Amdahl report.

- IBM, however, will not introduce a fighting machine but will respond in accordance with its long-term goal of preserving its growth and revenue objectives rather than destroying competition.

- The IBM response will be in two steps, both aimed at improving the price performance of the 158 and 168: the first step will be to upgrade the 158 and will occur by the end of 1976, followed by the upgrade of the 168. Both of these upgrades will be short-term responses to competitive pressures. The long-term (12-24 months) response will be to introduce a new generation of systems replacing the 158 and 168 mainframes.

A. SHORT TERM RESPONSES (NEXT 6-12 MONTHS)

- Lower the lease and purchase price of MOSFET memories.
- Integrate the data channels into the mainframe and reduce prices by 20%.
- Increase the size of the high speed buffer to increase throughput.
- Introduce a 158 APS with a 1.4 to 1.6 price performance improvement.
- Introduce a 158R* and 168R* and a 168 APS-s.
- Introduce a 2305 replacement to speed up the virtual memory operating system.
- Upgrade 3350 Disk Drive.

B. LONGER TERM RESPONSES (NEXT 12-24 MONTHS)

- Modify the architecture of the 158 and 168 by converting the mainframes into a functional multiprocessor containing an instruction processor, a service processor, a control processor, and one or more I/O processors, similar to the 115 and 125 design.

---

* Replacement

- Change channel interfaces and introduce security hardware codes between the I/O processors and the instruction processor. This would force competitive processors to emulate in software the IBM hardware features causing them to slow down.
- Put some of the OS software into hardware to make compatibility more difficult.
- Utilize the relocatable control storage and ROMs to put into hardware a variety of operating system kernels, applications, and language processors to increase throughput and reduce the operating system overhead.
- Increase number of instructions to facilitate VM.
- Replace its present components with higher density chips (16K memory chips are currently in production while 32K and 64K chips are in the prototype stage) to reduce component count, increase reliability, reduce costs and power requirement.
- Utilize a new memory technology such as bubble or CCD to reduce memory cost.

- Some of the above responses would be self-impacting; others may expose IBM to suits. IBM will announce those responses which, when simulated on its in-house system, will yield the optimum impact on competitors without significantly reducing revenues or increasing legal exposure. This judicious combination will also provide a solid base from which to launch the next generation of mainframes.

SPERRY UNIVAC

- The 1100 family is the base of all Univac large business systems. There are over 1200 medium and large 1100 Univac business systems installed throughout the world. The design evolved from the 1107 which was replaced by the single processor 1108 in 1964. Univac introduced a multiprocessor 1108 in 1965.

- The lower priced 1106 was introduced in 1969, and a multiprocessor 1106 was introduced in 1970.

- In 1970 the 1100 was announced. Each 1100 CPU has approximately 1.8 times the processing power of an 1108 CPU at 10% to 15% over the 1108 price.

- The 1106 II, announced in 1972, uses a faster memory than the previously announced 1106.

- The 1100/10, 1100/20 and 1100/40 were announced in 1975. They utilize MOS memories.

- In 1976 the 1100/80, the largest Univac computer, was introduced.

- The key technical features of the 1100 are:
    - Multiprocessing and hardware redundancy.
    - 2 levels of directly addressable memories (one very fast).
    - Instruction lookahead and concurrency (achieved by a four-deep instruction stock).
    - Decimal arithmetic and editing instructions.
    - Independent I/O processors.
    - Independent peripheral processors.

- Univac has made a major penetration into the following industries:
  - Federal government
  - State government
  - Service industries
  - Airlines telecommunications services
  - Educational institutions
  - Manufacturing
  - Printing and publishing
  - Savings and loans

- Univac depends heavily on the large business system market for its revenues. Currently, over 50% of Univac's revenues are derived from the 1100 product line, and the company believes that the LBM market will continue to grow at 15% per year for the next 5 years.

- The Univac 1100 operating system was originally designed for the 1108, but it has been extended to support the entire 1100 product line. Because of its evolutionary development, the Univac 1100 operating system is now mature and very comprehensive. It supports local and remote batch operations as well as interactive processing utilizing 1 or more processors.

- The use of re-entrant processing on Univac 1100 systems allows many jobs to share the same language processors. This technique saves memory and allows more jobs to operate simultaneously. Univac also provides:

  - a communications management system (CMS) to control messages

- a transaction interface package (T.P.) to control applications programs which are transaction-oriented.
- conversational time-sharing (CTS), providing users with an efficient man machine interface.
- a data management system (DMS) to provide data storage manipulation and retrieval and software.

- Univac is very aggressive in the large business market and is not reluctant to compete directly with IBM, especially in the multiprocessor environment where the mature Univac operating system and hardware approach give it an advantage over IBM.

- The 1100 product line is well balanced, and it can support both commercial and scientific computation, although the emphasis, because of the software and the limitations of the 36 bit word, is primarily on commercial computing.

- Univac management believes that users want availability, security, and ease of programming. They believe that a good data base management system is the key to a successful system. Univac prefers a hierarchical rather than a ring structure system.

- Five years from now, Univac management believes that the ease of interfacing a user to the computer will be an added concern. Utilizing a "dialog processor," they expect to be able to keep up with this user requirement.

- The industry is now close to reaching the theoretical speed performance limits of components. However, the utilization of components produced by using electronic beam lithography to further reduce component

size still has to be incorporated into the large business mainframes. his will include improved packaging techniques developed in-house.

- These advances will incorporate sub nanosecond ECL, very large scale integration, and high speed MOS memories.

- Univac believes that there will be one or more bubble memory devices announced before the end of this year.

- Univac indicates that it is on the verge of achieving all these breakthroughs. After they are incorporated, the next threshold may be the"Josephson Junction" (J.J.). There are many unanswered questions before a J.J. computer can be implemented. In fact, other techniques may surface or evolve in the next 10 years before a J.J. computer can be built; besides, increased speed is not the only key to a user-oriented system.

- The consensus at Univac is that IBM will continue to evolve its present product line and that the next announcement in the LBS arena will not be much more than another evolutionary step. Univac believes it can maintain a significant price/performance margin over IBM and continue to hold its profit margin over the next 3 to 5 years.

- 80% of Univac's R&D goes into software and new architecture where the company believes that new breakthroughs will occurr.

- One argument advanced by Univac against the multi-minicomputer concept is that memory is used ineffeciently in the minis, whereas a larger, faster system can share memories.

- There will be another substantial reduction in the cost per bit of disk storage as thin film heads and oriented media combine to make higher bit densities feasible.

- Univac has a Janus-like attitude on IBM compatibility. The 1100 line is definitely not compatible, and there are no plans to make it compatible. The 9000 line, whose upper end is similar in performance to the 370/158, is also a bridge to the Spectra 70 line. The 90/80 incorporates many 370 like features and can be made fully compatible with the 370. This is a controversial point within the company with opinions split between advocates of full compatibility and those who oppose it. The classical arguments are advanced by both sides to formulate their positions. The reality is that Univac has the best of both worlds, and if it can continue to stay on top of technology, it has the opportunity to develop the 1100 line, the 90/80 business, and its recently announced entry in the small business system market, the BC/7.

- The margins in the high end of the business will undoubtedly come down over the next 5 years. As the margins decrease, the level of R&D will undoubtedly decrease. This will hurt Univac, Burroughs, and Honeywell more than IBM, since IBM can distribute its R&D over a larger number of machines. Univac can, however, play both the "me too" and the "uniqueness" strategy, giving it perhaps an extra margin of flexibility.

- User reaction to Univac's performance is generally good. 5 out of 7 users interviewed liked their 1100 systems. 1 had problems with an Ampex add-in memory, and 1 indicated that, having been an IBM user,

he still longed for the "IBM support" but was willing to forego that because of Univac's lower price.

- The features most liked by users are:
    - Fail-soft and reliability (3)
    - Larger capacity, T/S and on-line programming capibility (2)
    - Reliability (1)

- The only complaint was the Ampex memory. No one else had any complaints.

- All but 2 systems were installed after direct head-on competition with IBM and others. The reasons for selecting Univac ranged from availability of good numerical control software, to benchmarks, to good communications capability. The key, in addition to solid hardware and software, seems to be a competent sales force not intimidated by IBM.

- Few Univac users seem even mildly concerned with IBM compatibility; one indicated that his tapes were IBM compatible and that that was sufficient. Another indicated that Univac had provided him with a special SDLC interface so that his system could communicate with IBM terminals.

- All users were planning to upgrade. Their plans ranged from a minimum 20% to a three-fold growth in mass storage, going from an equivalent of 50 3330s to 150 in 5 years. More communications, more on-line minis, more terminals and redundancy were other requirements. Thus, upgrades alone seem able to provide Univac with its planned growth in the large business system arena.

- The feeling among Univac users about the impact of micros and

minis is evenly split. 4 users believe that minis will take some of the load off the large system, whereas, 3 believe that although some work will be removed, there will be added applications which will continue to push up demand for the resources of the large system.

- The effect of lower communications on computing hasn't yet been thought through by most users. In fact, at this point, very few large users are using anything but the lines available from A.T.&T. for data communications.

- Univac captured all the accounts interviewed by running a successful benchmark and providing a better (30% average) price/performance than IBM.

Conclusion

- Univac has consolidated its product line and has successfully absorbed the RCA base. Its thrust into the small business processor market is rounding out a product line which was weak at the low ends. The ISS subsidiary is working on advanced disk drive designs to be announced within the next 12 months.

APPENDIX

NOTES ON "AN IBM VIEW OF THE FUTURE"

Talk given by Dave Norris, IBM,
to SHARE 48, Houston, on
Wednesday, March 9, 1977

Dave said his talk today will be regarding how IBM thinks about large systems. (That is the charter of his group).

He stressed (because IBM lawyers insisted) that his view was "an" IBM view, not "the" IBM view. (However, considering the forum (SHARE), and the fact that this talk was well publicized in advance, I feel the difference between "an" and "the" is probably close to non-existent.)

Dave Norris' group is in Poughkeepsie and is organizationally located here

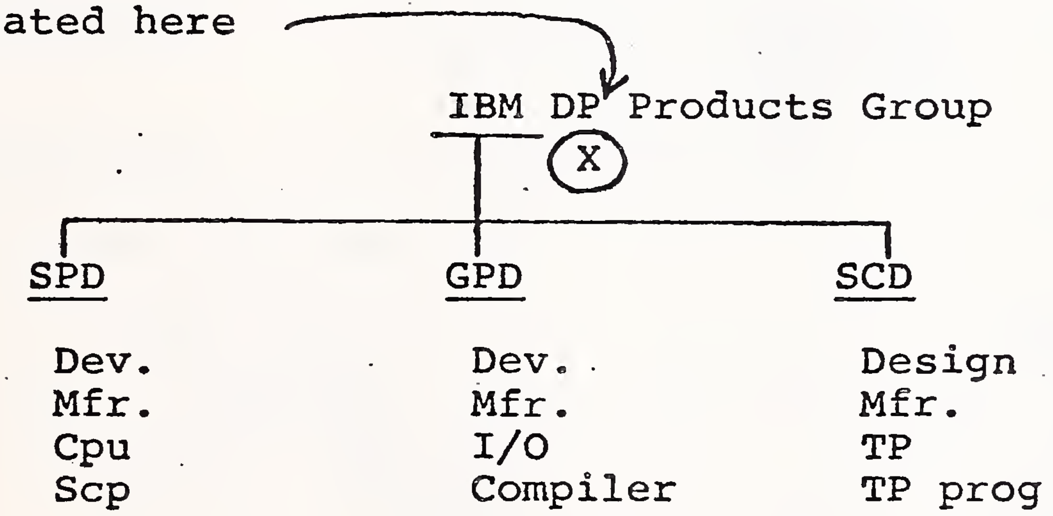

Norris's group takes a "systems" view (not a SPD, GPD or SCD view), and focuses on "futures" design for the next 5-10 years. Norris himself is particularly involved with (1) large processors, (2) system I/O, and (3) SCD TP.

Performance is the key consideration in the future. Performance can be obtained from

A. Technology
B. Machine Organization (e.g. System Balance)
C. Address/Access Space Management (e.g. Memory Hierarchies)

The old days of technology driving huge performance improvements are gone. IBM must obtain performance improvements from (B) and (C) above in the future.

MVS is an entirely different way of thinking about Data Base. It is only the tip of the iceberg.

"The real problem is in the central processing unit."

Look at evolution of computers to see what's going to happen next.

(1) Original Concept (back in "ancient" history)

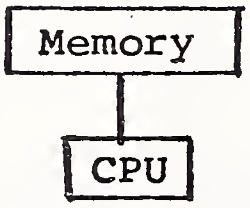

But memory was cheap and slow and CPU was fast and expensive. Found CPU was spending one half it's time doing nothing. So tried to remedy this.

Before going further, remember that any machine instruction is in three parts - - -

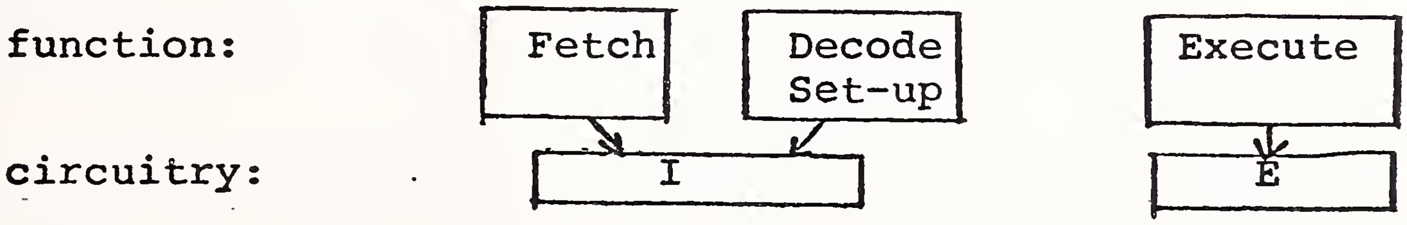

(2) Next step was

Model 75

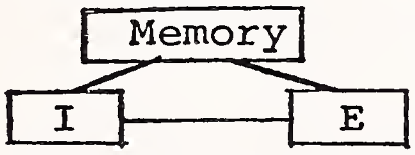

M91, M95

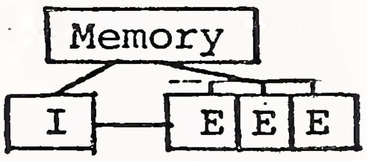

But only way to pay for these very expensive "E" boxes was to keep them busy all the time. But programmers don't program to keep "E" boxes busy, they program to develop applications. Thus, didn't sell many M91, M95.

Designed M195 to overcome these problems. But didn't sell many M195s either. That bothered IBM because they thought customers were saying "we want bigger, more powerful computers." IBM was building "bigger, more powerful computers," but they weren't selling.

So, in 1969-70 IBM decided to begin listening to customers more closely.

Task Force was created. Sent 30 people into the woods for a couple of months to determine what customers really do with these machines.

Task Force conclusions caused IBM to fundamentally change the way they thought about large systems.

Task Force report result: Distinction between

CAPACITY AND PERFORMANCE

CAPACITY = Ability to do a great number of things fairly quickly

PERFORMANCE = Ability to do something very, very quickly.

IBM needed to get into the high capacity (business) market, not the high performance (scientific) market.

Can't sell high performance machines to a high capacity market.

How build a high capacity machine? Have lots of "I" and "E" boxes to execute simultaneously.

Model 360/30

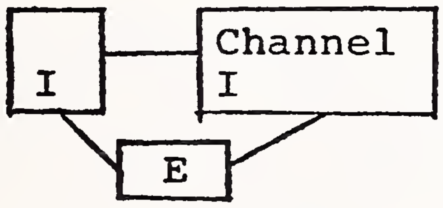

Above M30 approach is good for medium size machines, but for bigger machines need multiple "I" and multiple "E" approaches.

Want to spend some time discussing this approach because now we can make some good speculations on how systems will be built in the future.

First way to build them is - - -

Tightly Coupled (Share common control program)

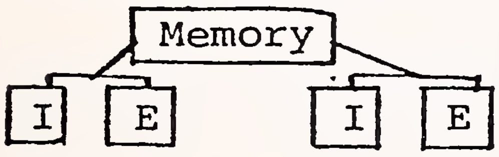

(1) Homogeneous Systems

MP65
AP168

MP158
MP168
AP158

(2) Non-Homogeneous i.e. Functional

e.g. 370/2938 Array
Processor

Problem = recovery

Microcode: People ask us "why is IBM so dedicated to Microcode?" One Reason - gives us tremendous flexibility in design of non-homogeneous systems.

Now we can speculate about how we might want to build systems in the future.

Take financial institutions for example: we could take a M168 and microcode it to be a 370. Then we could take a M148, bolt it on, but instead of microcoding it to be a 370, will microcode DL/I (or VSAM) right into the hardware. Take a third 370 system, bolt it on, microcode VTAM or BTAM or batch or whatever, right into the hardware. Take a fourth system, bolt it on, make it COBOL, DL/I or whatever you want.

Begin now, via microcode, to configure to exactly what kind of environment you want.

But the most important thing is that microcode can be loaded and changed dynamically. Many customers' workload changes dynamically during the day. For example, a bank does 70% of its daytime customer business during the two hour lunch period. That means that most of the terminals shut down and we begin to hammer to death the data base.

So, for a bank, we can take the VTAM engine and make the second load a DL/I microcode engine. So now have two engines hammering away at the data base. Take the 370 and make it a 370:80% of the time it will handle the VTAM work. At nighttime can make all the engines 370s or whatever you want.

Thus microcode gives excellent flexibility in configuring to an environment, and solves the recovery problem of non-homogeneous machines (if I lose my DL/I engine, simple load microcode into the other engine.)

But tightly coupled systems do have one problem in that they share one control program, and sometimes it fails. So, also developed - - -

Loosely Coupled Systems (not share a common control program)

Advantage: system failure on the left has no effect on system on the right.

INPUT

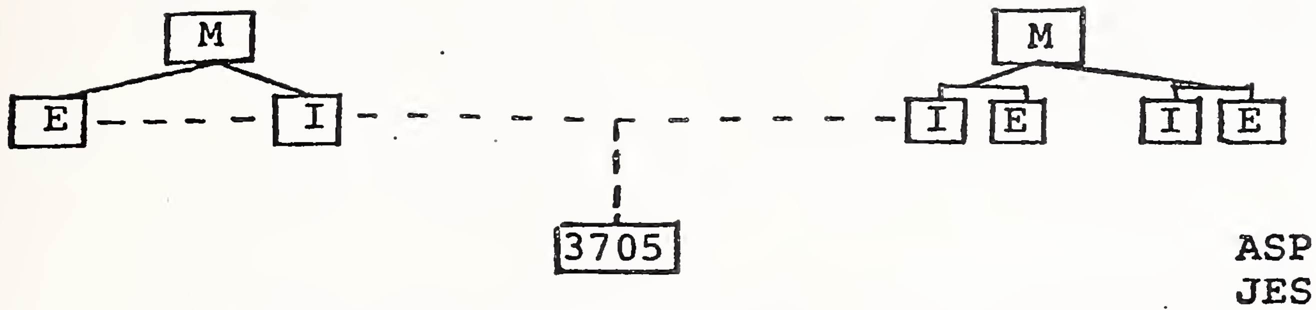

ASP
JES III
3704/3705
S/7, 3838

Faster unit processors -- not IBM's major future direction.

Makes more sense to couple systems together.

But once IBM made that decision, after 25 years of building faster and faster unit processors, it became clear that the software was going to have to change too.

VS1, SVS = batch software, we're not talking about batch environments. This is "on line data base type stuff" we're dealing with now.

On-line = a whole different philosophy. Not know what my next transaction is. Demand schedules.

Thus, it became clear to IBM they would have to completely rethink the way they did software, not only for coupled systems, but for the whole new data base environment. Its what I call - -

Addressing and Memory Hierarchies

Next, want to show you what caused us to come to the conclusion that the old software wouldn't suffice.

Let's go back in history again.

7080 (late 1950s large addressable memory)

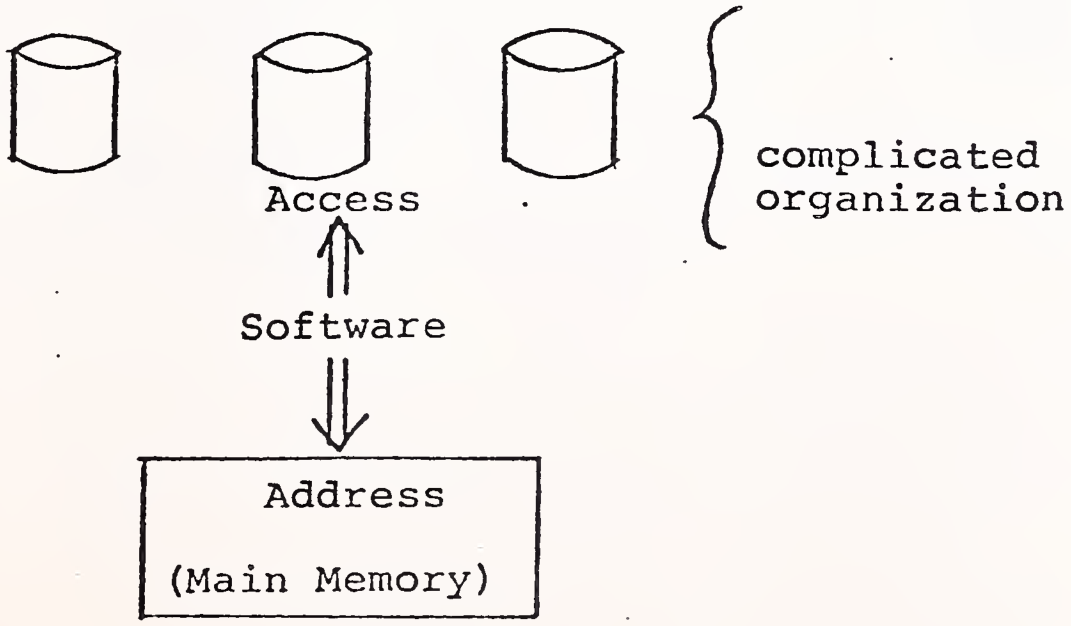

INPUT

(1) Keep in mind that data on the access devices (tapes, disks, drums) is absolutely useless. Can't add data or do anything. Must take an enormous amount of software to locate that data, retrieve it, and get into the place where something can be done, i.e. addressable memory.

(2) Device - independence. If change disks - impacts software.

(3) Disk organization, etc. is very complicated. In addressable memory, not so complicated.

In other words, access devices are "disaster areas". Only place we can do productive work is down in addressable memory.

In 360, tried to improve the situation by coming up with a hierarchy of memory.

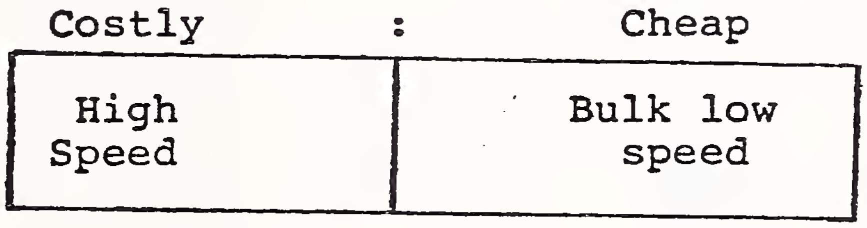

Memory

This was our first attempt to try to reduce need for that complicated access software. But few people used the internal high-to-low speed channel.

So, decided hierarchies were good ideas, we just implemented it wrong. So our next try - - was "cache."

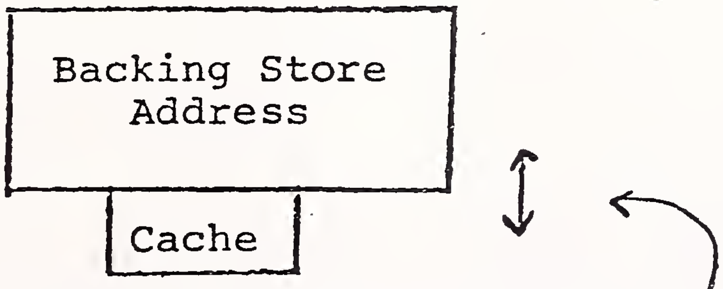

4,000,000 bytes
M85, M195
M155, M165

this boundary managed by hardware

But, meanwhile mechanical devices, (e.g. disks) were still getting bigger and cheaper, faster than electronic devices.

World of data base was now becoming more prevalent. But it takes customer much money and time to implement a data base system.

We felt that bubble memories (billions of bytes) were coming in late 1970s, early 1980s. But back in 1972, couldn't wait for that. So decided to solve the problem and add a third hierarchy of memory. Each user could have 16 million bytes of memory. But since couldn't afford real memory, decided to implement with old drums and disks, and use a combination of hardware and software. (Hardware was the relocate feature).

So IBM microcoded the M148, M158 and M168 to address 16 million bytes per user. But since don't really have 16 million bytes of real memory, had to develop the Paging Supervisor to bring data from the disk into available memory.

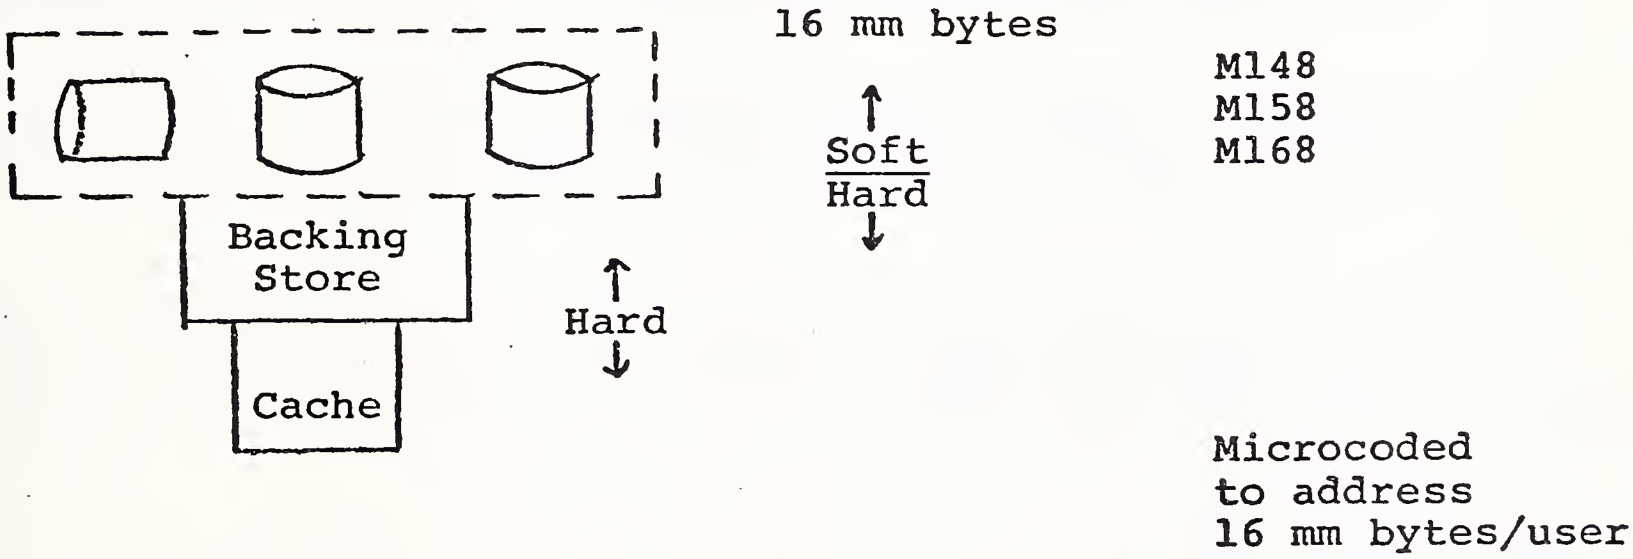

People ask IBM "why did you do it that way?" We did it for two reasons:

When new bubble or thin film memories come along, two things will happen:

(a) Since the microcode for addressing 16 million bytes is already in the computer, the user has already programmed for it. Thus it will protect the programming investment.

(b) When we add a function, systems software goes away. Thus as we get bigger and bigger memories, "Virtual" approaches "Real", and we page less and less. Someday it will completely go away. So the systems software begins to disappear.

Next, we stepped back and said "if we can be so clever down in the memory area using software/hardware combinations, why can't we do the same up in the real disk world -- the "disaster area."

We tried that. We introduced "Virtual I/O under MVS" to try to do that. The concept of virtual I/O is to take one of the disks up in the "disaster area" and, unknown to the user, move it down to one of the paging disks. The user still programs like he always did. But rather than using the cumbersome software up above, he is using the microcode and is very efficiently moving the data back and forth. This approach can eliminate up to 50% of the instructions necessary to transfer that data.

But now we had a dilemma. Don't want to move things into microcode that are not stable. And things like DL/I and VSAM are not stable.

We implemented Virtual I/O for sequential access because that access method is stable.

But it is clear what we want to do. As the other more complex data base access methods stablize, we want to move them into Virtual I/O and microcode.

In view of the tremendous explosion in the storage of data on tape and DASD, we felt it would be useful to try to use some of the clever virtual techniques up there in the "disaster" area.

So, we introduced the 3850 MSS. I would like today to use the MSS to show you what MVS is all about, to show you what we are really trying to do. So far what I've talked about is not radically different.

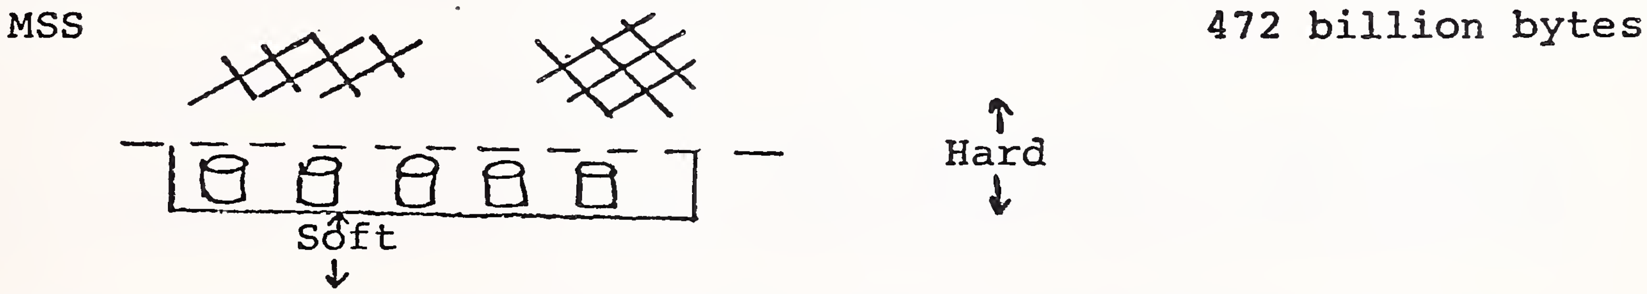

We fooled the software into thinking we had an almost unlimited number of disks.

Look at the capacity designed into MVS -- 472 billion bytes! (NY Public Library book collection is only 2/3s of that!)

Now look at MVS. It has the capacity to directly address 512 billion bytes of data.

So now we can begin to see for the first time what we are really up to with MVS.

We are saying --- we don't really believe data bases belong out on those "external" devices. We think data bases ought to be down in main memory.

So we designed a hardware and software combination that will begin to allow us to migrate the data base down here where it really belongs -- in the main memory.

It has been only recently that programming announcements have come out that show you what we are really trying to do.

One in particular is IMS FAST PATH, which was announced Sept./ Oct. 1976. The importance of IMS FAST PATH is a feature called "IMS Main Storage Data Base." The IMS Main Storage Data Base is not up on the disks, but is down in memory where we really think it belongs.

What we are saying is that IMS Main Storage Data Base is not important for what it is today, because right now it has a lot of restrictions.

What is important is what IMS Main Storage Data Base is trying to tell us -- namely, that there are announcements coming out which tell us what MVS is all about.

Now we can begin to very reasonably speculate what these systems will be looking like in 1985-1988.

Data Base belongs in main memory.

Now how do we get CAPACITY out of that data base? -- We surround it with coupled (tightly, loosely, homogeneous or non-homogeneous) systems.

What we're trying to say is that we have seen two computer revolutions in the early 1970s.

A) The first revolution is "Technology." Many people misunderstand it. They think it is "speed." But it isn't, it is "cost." A hand held calculator of 1977 is no faster than a hand held calculator of ten years ago, but it costs $8 and has more functions than the one of ten years ago which cost $800.

We don't think it makes sense to build a blazing fast, stupid box. It makes more sense to build a very smart, very fast box. In other words --

coupled systems with lots of microcode, and with the ability to take that moving line of hardware/software interface, into the hardware.

B) The second revolution: software is driving the computer industry, not hardware.

For 25 years technology has driven the computer industry. Now it has completely flipped around. Thus, we had to design a hardware/software architecture that would allow us to begin programming for a 5, 10, 15 year time frame. We don't have the technology to do it, so we had to begin using virtual.

Thus we don't care what direction technology takes just so we don't have to throw away millions of dollars of software investment. That was the really big switch.

In closing, remember that Von Neumann is generally considered to be the father of the modern computer. His fundamental contribution was that there should be no difference between programs and data. Every computer vendor claims to be making Von Neumann machines. But none of us have - look where we have been putting programs (down in memory), and where we have been putting data (up in "disaster area" devices like tapes, disks, etc.). We can't execute anything when it is out on a device.

Where we think the future is, what VS is all about, is IBM's attempt to build a Von Neumann machine. We think he was right. There should be no difference between programs and data. They both ought to be able to reside down in main memory, where they can be effectively used.

That's the "IBM View of the Future."

[illegible]

USER QUESTIONNAIRE

SYSTEM CONFIGURATION

1. Do you have plans to expand?

		Present System	Planned Additions/Changes
2. System mfg. and model			
Memory size		K	K
No. of channels			
Type of peripherals			
Disk	IBM		
	Other		
Tape	IBM		
	Other		
Printers	IBM		
	Other		
RJE	IBM		
	Other		
Terminals	IBM		
	Other		
Operating System			

Installed When? ____________

3. Other computers in installation? ________No ________Yes

Manufacturer	Model	Core Size
		K
		K
		K
		K

4. How long do you plan to keep your present mainframe sytem?_______years.

5. Have you any concern that your present commitment will prevent you from taking advantage of future technological improvements?

Hardware?______________________________

____

Software?______________________________

____

Communications?______________________________

____

6. What types of changes would cause you to look at other manufactuers mainframes?

Price?______________ What differential?______________

Ease of use?______________________________

Other?______________________________

IBM compatibility?______________________________

SATISFACTION RATING

7. Is the sytem performing to your expectations? (Downtime, throughput, etc.?)______________________________

____

____

8. Was the conversion smooth?______________________________

____

____

9. What do you most like about the new system?______________________________

____

____

10. Any negative opinions about present supplier support?

Technical?

Software?

MARKETING

11. What did competition do to keep or capture the account?

12. What were the strong points of marketing and marketing support personnel? (Technical skills, understanding of user, economic analysis, etc.)

FUTURE CHANGES

A. What changes do you see in your customers over the next 5 years?

B. What will the effect of Distributed Processing be on your Large Business System?

C. What will the effect of mini's and micro's be on your LBS?

D. What will the effect of lower cost data communications be on your LBS?

E. What will the effect of word procession be on your LBS?

SELECTION PROCESS

13. Why did you select your present MF supplier?

14. How many people involved in the study?

Titles?

15. How many involved in the decision? ____________

Titles?____

____

____

____

16. Was decision unanimous? _____Yes _____No

17. Other

INPUT

VENDOR QUESTIONNAIRE

1. How big is the large business system (LBS) market today? (Over $2 million systems)

	Total Market		Your share	
	Units	$	Units	$
1976 Shipments				
1976 Installed Base				

2. How big do you expect the LBS market to be in 5 years?

	Total Market		Your Share	
	Units	$	Units	$
1981 Shipments				
1981 Installed Base				

Why? ____

____

____

____

3. What is your perception of the major concerns of the users in the LBS market today?

4. Five years from now?

5. What Component Technology/Software/Architecture development do you see impacting the LBS market in 5 years?

   A. Component Technology______________________________________

   ____

B. Software

C. Architecture

D. Multi-level Logic

6. What Price/Performance improvement do you see over the next 5 years?

Why?

7. What will be the impact on the LBS market in the next 5 years of the following?

A. Distributed Processing

B. Distributed Data Bases

C. Mini's and Micro's

D. Storage Costs Mainframe

E. Auxiliary Memory

F. Word Processing & Office Functions

G. IBM Capability e.g. AMDAHL

8. How do you sell the LBS market?

   A. Dedicated Sales Force?

   B. Industry Oriented?

   C. Targeted?

9. Shipment History

	Model		Model		Model		Model		
Year	Units	$	Units	$	Units	$	Units	$	Total
1971									
1972									
1973									
1974									
1975									
1976									